LEVEL II

AD A090238

DTIC
ELECTE
OCT 1 4 1980
S C

UNIVERSITY OF MARYLAND
COMPUTER SCIENCE CENTER

COLLEGE PARK, MARYLAND
20742

DDC FILE COPY

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

80 10 3 110

TR-855
DAAG-53-76C-0138

January, 1980

EXTRACTING LINEAR FEATURES FROM IMAGES USING PYRAMIDS

Michael Shneier
Computer Vision Laboratory
Computer Science Center
University of Maryland
College Park, MD 20742

ABSTRACT

A method is described of extracting linear features from images. The approach is to construct a series of lower-resolution versions of the original image (a pyramid), and to look for lines in these images. A line in a low-resolution image corresponds to a thicker linear feature in a high-resolution image. The position and extent of this linear feature is calculated from the low-resolution image, and a threshold is found which, when applied in the neighborhood of the feature in the high-resolution image, segments the linear feature from its background. Advantages of the method are that only the parts of the image in the neighborhood of linear features need be thresholded, and that different thresholds may be used to extract the various linear features in the image.

The support of the Defense Advanced Research Projects Agency and the U.S. Army Night Vision Laboratory under Contract DAAG-53-76C-0138 (DARPA Order 3206) is gratefully acknowledged, as is the help of Kathryn Riley in preparing this paper. The author is lso grateful to Shmuel Peleg and Les Kitchen who supplied the line-enhancement programs.

1. Introduction

Linear features in aerial photographs are usually strong evidence of human activity, and can also be useful for navigation purposes. Linear features also occur in other kinds of images, such as bubble chamber photographs and industrial radiographs. This paper presents a method of finding those parts of an image that correspond to linear features, and of segmenting these features from their background.

The method involves constructing a series of successively lower-resolution images (a pyramid) from the original image (Hanson and Riseman, 1978, Tamimoto, 1978, Uhr, 1978). A line-detector is run over each of the low-resolution images, and its output is used to calculate regions in the original image that correspond to linear features. Lines detected in successively lower resolution images correspond to successively thicker linear features in the original image. Properties of the regions in the low-resolution image that give rise to line-detector responses are used to calculate a threshold for extracting the linear features in the original image. This work is an extension of an earlier system for detecting blob-like objects using pyramids (Shneier, 1979).

2. The Algorithm

The algorithm performs two tasks. First, it finds the parts of the image that correspond to linear features, and then it calculates a threshold to extract these features. Both tasks make use of the pyramid structure.

1. If the whole pyramid has been constructed, stop. Otherwise, read in the previous pyramid level (the picture, if this is the first iteration).
2. Build a new level.
3. Apply a line-detector to the new level (see below).
(4. Run a line enhancement algorithm on the new level).
5. Group the responses from steps 3 and 4 into lines (see below).
6. For each line
   a. calculate a threshold (see below)
   b. apply the threshold to the region in the original image corresponding to the line and write the results to the output picture.
7. Go to 1.

A pyramid level is constructed from its predecessor by replacing each two-by-two neighborhood of points by their median value. (It is also possible to use the average value, but this results in more blurring.) Thus, each successive level is one-quarter the size of the previous level.

The line detector that is run over the new low-resolution image was designed to take the neighborhood size into account. It has to be able to deal with lines that are two pixels wide, because it is possible that such lines might otherwise "fall between the cracks" when the next level is constructed. That is, one part of the line might become part of a different neighborhood than the other part, and the resulting points might no longer be detectable as belonging to a line. For horizontal and vertical lines, a three by five pixel mask is used (Figure 7), while for diagonal lines, a five by five mask is used (Figure 2).

Each mask is applied at every point in the image, and the final value for the point is calculated as a function of the responses. Thus, if two masks respond (for example, the horizontal and 45° masks) an interpolated value is calculated as a weighted average of the responses, usually resulting in a direction somewhere between those of the two masks (e.g. 30 degrees). The responses are all reduced modulo 180° to enable simpler tests for similarity to be made in the enhancement phase described below. Each point has a magnitude and direction associated with it. Figure 3 shows an example of the output of the line detector applied to a picture of part of an airport (Cumberland Municipal Airport in Maryland).

An optional, but usually advisable, step in the algorithm is the application of a line enhancement program to the raw

line-detector output. The purpose of this step is to remove noise from the line-detector output, and to adjust the directions of neighboring points to make them more compatible. The enhancement program that was used was adapted from that described by Peleg (1978).

For each neighbor Q of a point P, let M_Q and D_Q be the magnitude and direction of the line response at Q; let D_P be the direction of the response at P, and let D be the direction of the line joining P to Q. Then the degree to which the response at Q supports that at P is

$$S_Q \equiv M_Q \cos(D_Q - D)\cos(D_P - D)$$

Note that this is a maximum value when D_P and D_Q are both collinear with D. Using this function, new estimates of the magnitude and direction of the edge at P can be computed as follows.

a. The new magnitude is proportional to ΣS_Q, where the sum is taken over all the neighbors Q of P.

b. The new direction is basically the direction to that neighbor Q_i for which S_Q is greatest, but modified as follows: Let the direction to Q_i be 45i degrees. Then the responses at Q_{i+1} and Q_{i+5} tend to bias the direction toward an angle greater than 45i degrees, while the responses at Q_{i-1} and Q_{i+3} tend to bias it towards a smaller angle. The final direction estimate is thus

$$45i° + \frac{S_{i+1} + S_{i+5}}{5} 22\tfrac{1}{2}° - \frac{S_{i-1} + S_{i+3}}{5} 22\tfrac{1}{2}°$$

where S_j is short for S_{Q_j} and $S = S_{i-1}+S_i+S_{i+1}+S_{i+3}+S_{i+4}+S_{i+5}$.

c. The magnitude of the response is set to zero unless one of the following conditions is met: Let Q_1, Q_2, Q_3 be the points at which the response magnitude is greatest, second greatest, and third greatest, respectively. Then Q_1 must be opposite Q_2, Q_2 must be opposite Q_3, or Q_3 must be opposite Q_1. These conditions correspond to the fact that for a straight line, the response magnitude should be maximum for a pair of opposite directions.

Peleg recommends iterating the above procedure several times. For the current method, however, a single application was used. As a result, the main advantage was the noise-cleaning effect of step c. Figure 3c shows the result of running the enhancement process on the output in Figure 3b.

The next step is to group the line-response points into line segments. This is done by means of a stepwise clustering process. First, the points are clustered into groups with similar direction, and then, within each group, a further subdivision is made on the basis of the separation between the points. Separation is measured by rotating the points so that their direction is vertical, and measuring the distance between them. This distance must be less than a fixed threshold for the points to belong to the same line. Because there is a fairly large amount of error introduced by the rotation, the rotated points are sorted in

order of increasing distance from the origin, and each successive point is added to a line segment if its distance from the previous point is less than the threshold. This provides a reasonable amount of laxity in the grouping process. If there are too few points in the line (currently, if there is only a single point) then the line is discarded. Similarly, if the length of the feature in the image to which the line corresponds is much greater than the number of points found on the line (currently, if there are points corresponding to less than half the length of the line) the line is assumed to be noise, and is also discarded. This prevents widely separated, isolated points that coincidentally line up from being considered as valid features.

Each line in the low-resolution image corresponds to an elongated region in the original image. The position and extent of this region must be calculated, as well as a threshold for extracting the corresponding linear feature.

At first sight, it appears easy to calculate the position and extent of a region from the position and level in the pyramid of the corresponding line. The endpoints of the line can be found and mapped into the ends of a region whose extent depends on the level of the line in the pyramid. Unfortunately, there is no guarantee that the endpoints of the line are reliable estimates of the real position of the line. Due to errors in the line-detector and enhancer (to which the endpoints are particularly

sensitive) and to peculiarities introduced by the fixed coordinate grid and the pyramid process, any individual point may lie significantly distant from the real center of the linear feature. As a result, the region calculated may only partially overlap the feature it represents.

An alternative approach is to calculate the average slope of the line from the slopes of the points, and to calculate a region based on passing a line with this slope through the centroid of the points. This approach has the disadvantage that it is hard to calculate the width and length of the region reliably.

The method that was employed in producing the examples relies on local rather than global properties of the points that make up a line. Points are treated pairwise along the line, and a series of regions is calculated, one for each successive pair, based on the coordinates of the points. This insures that the local slope is respected, and also insures continuity. It has the disadvantages of requiring more computation and of sometimes giving a slightly jagged look to the output.

A threshold is calculated for each line segment based on gray-level values in the low-resolution image. The threshold is the average of a central value, derived from pixels on the line, and an outer value, derived from pixels on either side of the line. Specifically, the central value is the gray-level of the point in the low-resolution image that gave rise to the

line-detector response. The outer value is the average of the gray-level values at two points on either side of the line point. These two points are calculated based on the direction of the line point, and are as close to the normal to this direction as is allowed by the co-ordinate grid. The points that are chosen are two pixels away from the line point in the case of lines that are approximately vertical or horizontal, and one diagonal pixel away in other cases. The threshold value for a line segment is the average of the threshold values of its endpoints. The threshold is applied to the points in the original image in the region calculated for the segment, and the results are written to the output image. When all lines have been processed, a new pyramid level is constructed, and the process is repeated.

3. Examples

Figure 4a shows a small part of Frederick Municipal Airport in Maryland. Figure 4b shows the parts of the image that were chosen to be thresholded, and Figure 4c shows the thresholded output. It is clear that even using local calculations to obtain the region corresponding to a line does not always give ideal results.

Figure 5 shows three other airport scenes, including the final result of the process applied to Cumberland Airport (shown in Figure 3). Most of the linear features are extracted. It should be noted that thresholding the entire image using a single threshold value is not acceptable for these images. The results of such a process either do not contain all the features, or contain large areas of the background as well (Figure 6).

Figure 7a shows part of a cloverleaf intersection. There is a substantial amount of curvature in this feature, and, as Figure 7b illustrates, the straight-line favoring process almost completely ignores the feature, even though a conventional thresholding technique would not be able to discriminate against the curvature. Figure 7c shows the result of applying the process without the enhancement step. Now much more of the cloverleaf is visible, but at the cost of less detail in the rest of the image.

4. Discussion

The examples in the previous section show that the process described is selective for linear features. This kind of selectivity can be useful in its own right, for example, when looking for roads or runways in images. In many applications such a process might also be useful as a first step in a more complex procedure. Suppose, for example, that an expensive operation is to be performed on an image, but need only be applied in restricted regions (e.g. counting cars on a road (Quam, 1978)). The procedure described in this report can then be applied to find relevant areas in which to apply the operation. By suitably changing the feature detector, it would be possible to select for arbitrary kinds of regions.

The method is also potentially useful in restricting the search space for other procedures. Another approach to linear feature extraction that is being actively studied is to use an edge-extraction technique to find the sides of linear features, and a clustering technique to pair up the edges that are extracted (Scher *et al*., 1979, Nevatia and Babu, 1978). This technique leads to difficulties because of fragmented edges, and the problem of deciding which edges should be paired to form the linear features. These difficulties could be reduced by first finding approximations to the features using the method described in this paper, and using the approximate results to restrict the amount of search for pairs.

One of the problems that became apparent in applying this method was the large decrease in resolution between pyramid levels. Each level is one-fourth the size of its predecessor, and features tend to blur significantly between levels. A further difficulty arises from the discrete neighborhoods used in calculating pyramid points. This tends to further blur the image by splitting features and merging parts of features with the background. A way of avoiding this problem is to use a different rule for constructing new pyramid levels. Burt (1980) gives rules for constructing pyramids that taper less sharply, and whose levels are constructed using overlapping neighborhoods. It is intended to implement some of his methods and evaluate their performance in this domain.

5. Conclusions

A method of segmenting parts of an image with specific properties has been presented. The process is able to find the linear features in an image, and segment these from the background even in the presence of non-linear features with similar gray-level characteristics.

The method makes use of a feature-detector (a line finder) and a pyramid of low-resolution images. The feature-detector finds those parts of the image that have the desired properties. It is then possible to apply further processes only in these restricted parts of the image. This results both in cheaper procedures and cleaner output.

References

1. P. J. Burt, Fast hierarchical convolutions with Gauss-like kernels (to appear).

2. A. R. Hanson and E. M. Riseman, Segmentation of natural scenes. In Computer Vision Systems (Hanson and Riseman, eds.), Academic Press, New York, pp. 129-163, 1978.

3. R. Nevatia and K. R. Babu, Linear feature extraction. In Proceedings DARPA Image Understanding Workshop, Pittsburgh, PA, November 1978, pp. 73-78.

4. S. Peleg and A. Rosenfeld, Straight line enhancement and mapping. Computer Science TR-694, University of Maryland, College Park, MD, September 1978.

5. L. Quam, Road tracking and anomaly detection. In Proceedings DARPA Image Understanding Workshop, MIT, Cambridge, MA, May 1978, pp. 51-55.

6. A. Scher, M. Shneier, and A. Rosenfeld, A method for finding pairs of anti-parallel straight lines. Computer Science TR-845, University of Maryland, College Park, MD, December 1979.

7. M. Shneier, Using pyramids to define local thresholds for blob detection. In Proceedings DARPA Image Understanding Workshop, USC, Los Angeles, CA, November 1979, pp. 31-35.

8. S. L. Tanimoto, Regular hierarchical image and processing structures in machine vision. In Computer Vision Systems (Hanson and Riseman, eds.), Academic Press, New York, pp. 165-176, 1978.

9. L. Uhr, "Recognition cones", and some test results; the imminent arrival of well-structured parallel-serial computers; positions, and positions on positions. Ibid., pp. 363-378.

d		a		g
e		b		h
f		c		i

Figure 1. Mask for vertical line detection. Point b is accepted as a vertical line point if a > d and a > g, b > e and b > h, c > f and c > i. The horizontal line detector is analogous.

		d		
	e		a	
f		b		g
	c		h	
		i		

Figure 2. Mask for diagonal line detection. For a 45° line, point b is accepted if a > d and a > g, b > e and b > h, c > f and c > i. The 135° detector is analogous.

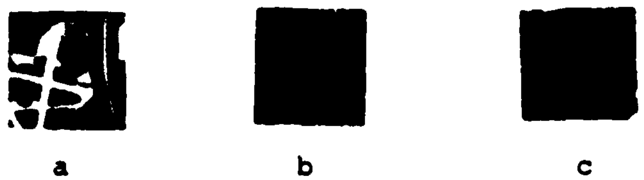

Figure 3. a. First pyramid level (Cumberland Municipal Airport). b. Thresholded line-detector output. c. Thresholded enhanced output.

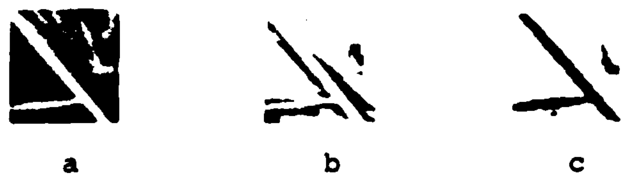

Figure 4. a. Part of Frederick Airport. b. The regions chosen for thresholding. c. The thresholded output.

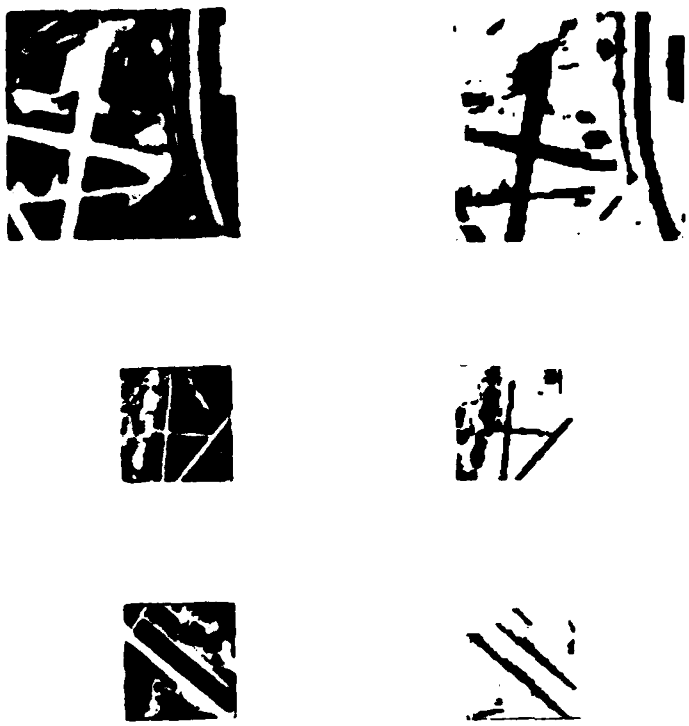

Figure 5. Three airfields and the results of applying the linear-feature detector to them.

a b

Figure 6. a. A global threshold producing results similar to those in Figure 5 is obtainable only at the cost of missing a vertical linear feature. b. When all the linear features are extracted, so is a large amount of the background.

a b c

Figure 7. a. A cloverleaf intersection. b. The result of applying the detector to (a). c. The result of applying the detector without line enhancement.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER

2. GOVT ACCESSION NO. AD-A090238

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle)

6 EXTRACTING LINEAR FEATURES FROM IMAGES USING PYRAMIDS.

5. TYPE OF REPORT & PERIOD COVERED

1 Technical rept.

6. PERFORMING ORG. REPORT NUMBER

14 TR-855

7. AUTHOR(s)

10 Michael/Shneier

8. CONTRACT OR GRANT NUMBER(s)

15 DAAG 53-76C-0138

V/DARPA Order-3206

9. PERFORMING ORGANIZATION NAME AND ADDRESS

Computer Vision Laboratory, Computer Science Center, University of Maryland, College Park, MD 20742

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS

11. CONTROLLING OFFICE NAME AND ADDRESS

U.S. Army Night Vision Laboratory
Ft. Belvoir, VA 22060

12. REPORT DATE

11 Jan[redacted]80

13. NUMBER OF PAGES

18

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)

12 21

15. SECURITY CLASS. (of this report)

Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Image processing
Pattern recognition
Linear features
Segmentation
Pyramids

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A method is described of extracting linear features from images. The approach is to construct a series of lower-resolution versions of the original image (a pyramid), and to look for lines in these images. A line in a low-resolution image corresponds to a thicker linear feature in a high-resolution image. The position and extent of this linear feature is calculated from the low-resolution image, and a threshold is found which, when applied in the neighborhood of the feature in the high-resolution image, segments the linear

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

411074 Jm

feature from its background. Advantages of the method are that only the parts of the image in the neighborhood of linear features need be thresholded, and that different thresholds may be used to extract the various linear features in the image.